

ETHERNET-TO-ATM INTERWORKING TECHNIQUE

Technical Field

[0001] This invention relates to a technique for interworking different types of data networks that have different protocols.

Background Art

[0002] Presently, communication service providers, such as AT&T, offer high-speed Asynchronous Transport Mode (ATM) Virtual Private Network (VPN) service to customers. Each ATM-based VPN customer utilizes one or more Permanent Virtual Circuits (PVCs) to route data among different locations (endpoints), each typically located at a separate edge of an ATM network. In practice, traffic originating at an endpoint passes to an edge device on the ATM network for transmission to the network core, which in turn, transmits such traffic to an edge device serving the destination endpoint. While the edge devices may run one or more different protocols, including ATM or Frame Relay, the network core typically utilizes the ATM protocol. In this environment, ATM PVCs constitute a point-to-point network topology.

[0003] Currently, there exists a large embedded base of Ethernet Local Area Networks (LANs). Advances in Ethernet technology have lead to the development of Metropolitan Area Networks (MANs) that afford access to the Internet and some limited access to VPNs. Ethernet-based MANs offer significant cost advantages on a per port basis, as compared to Frame Relay and ATM networks. Many VPN customers would like the opportunity to use an Ethernet-based MAN to access their ATM-based VPNs but have not had the ability to do so because of interworking issues. The protocol associated with Ethernet is different than that associated with ATM. Ethernet is a broadcast protocol within level 2 (the data link layer) of the well-known 7-layer OSI model, whereas ATM and Frame Relay is a point-to-point circuit-type protocol within level 2. Ethernet is designated as a broadcast protocol within level 2 because information in an Ethernet network travels in both directions and passes by all devices on the path. A device that

recognizes the information intended for itself (as opposed to another device) will pull the information from the network.

[0004] Thus, in the past, a customer seeking to use an Ethernet-based MAN to route traffic to a VPN served by an ATM network had to worry about both Ethernet and ATM protocols. Interconnecting these two protocols typically required a high level device like a router.

[0005] Thus, there is a need for an interworking technique that enables a customer on a first network, such as an Ethernet MAN, for example, to send information to an end point on a second network, such as an endpoint on an ATM network, without any concern as to the protocol of the network serving that endpoint. Furthermore, this technique must be able to interwork between a broadcast domain and a point-to-point circuit-based-domain.

Brief Summary of the Invention

[0006] In accordance with one aspect of the invention, there is provided a technique for sending information from a source to a destination. The information is embodied in a payload of at least one frame sent from the source served by a first network connected by a second network to the destination. In accordance with the method, an interworking facility receives frames that are destined for the second network, each such frame destined for the second network including not only the payload, but also a destination address indicative of the endpoint in the second network destined to receive the information in the payload. The destination address is obtained by initially resolving the destinations available to the source, including those available through the second network. In practice, the interworking facility establishes a set of pseudo addresses in a format compatible with the first network that correspond to destinations in the second network so that the source can address an information frame using its own protocol for a destination that actually lies in the second network without concerning itself with the protocol employed in the second network. In the case where the first information frame comes from a source in an Ethernet-based network, the first information frame will have a Virtual Local Area Network (VLAN) tag associated with the address of the destination. On the other hand, if the information frame comes from a source in an ATM network, the

frame will include a VPN Virtual Circuit Identifier (VCI), herein after referred to as a Permanent Virtual Circuit (PVC) that corresponds to the address of (e.g., the network path to) the destination in a format compatible with the ATM network, even though the destination lies in another network having a different protocol.

[0007] Upon receipt of the first information frame at the interworking facility, the facility forms a second frame compatible with the second network, the second frame including the payload. The destination address of the first frame is mapped to a second destination address compatible with the second network. Thus, for example, the VLAN tag in an originating Ethernet frame is mapped to a VPN PVC for an ATM frame and vice versa. Mapping the destination address from a format compatible with the first information frame to a format compatible with the second information frame allows routing of the second frame, including the information embodied in its payload, to the destination.

[0008] In accordance with another aspect of the invention, there is provided a technique for accomplishing address resolution for a source served by a first network to enable it to establish at least one available destination for receiving data, including a destination available through a second network having a protocol different from the first network. To accomplish such address resolution, the source broadcasts to an interworking facility an Address Resolution Protocol (ARP) polling request for the purpose of identifying each available destination, and in particular, an identifying address for that destination. Upon receipt of the ARP polling request, the interworking facility matches an identification tag in the request (e.g., the Virtual LAN Identifier tag for an ARP polling request from an Ethernet-based source) to an address for the destination in the second network (e.g., a path identifier, such as a Permanent Virtual Circuit (PVC) for an ATM-based destination). The interworking facility encodes the ARP polling request into a format compatible with the second network and transmits that request to the destination. In response, the destination replies with its address to the interworking facility that translates the destination-identifying address into a format compatible with the first network for transmission therethrough to the source. For example, in the case of an Ethernet-based source, the interworking facility encodes the ARP reply from the destination with a Media Access Control (MAC) layer address from a pool of local

addresses associated with the Ethernet-based source for transmission thereto. Under such circumstances, the ARP reply would also include the IP address of the ATM-based destination and the VLAN of the Ethernet-based source. Upon receipt of the encoded destination-identifying address, the source can thus identify the destination and send information thereto such that the destination appears as to the source as an end point on the first network. In actuality, the interworking facility acts as a proxy in the exchange between the source and destination.

Brief Description of the Drawings

[0009] FIGURE 1 depicts a block schematic diagram of a network architecture for practicing the interworking method in accordance with the present principles; and

[0010] FIGURE 2 depicts a block schematic diagram of a customer router and an Ethernet Interworking Switch, comprising part of the network architecture of FIG.1, and the manner in which the EIWS acts as proxy between the Ethernet domain and the ATM domain; and

[0011] FIGURE 3 depicts a block schematic diagram of an Ethernet Interworking Switch, comprising part of the network architecture of FIG. 1, and the manner in which the switch maps destination addresses from one format to another.

Detailed Description

[0011] FIGURE 1 depicts a block schematic diagram of a network architecture 10 in accordance with a preferred embodiment of the invention for interworking a source and destination that lie in first and second networks having different protocols to allow the source to send data using its own protocol by first establishing for the source a set of address in a format compatible with the for destinations that lie in second network, and thereafter having an interworking facility act as a proxy between networks. In the embodiment of FIG. 1, the source 12 comprises a first router and the destination, comprises one of routers 14, 16 and 18. In the illustrated embodiment, the first router 12 routes traffic, in the form of Ethernet-formatted information frames 20 (only one of

which is shown), onto an Ethernet-based Metropolitan Area Network (MAN) 22 which comprises a first network. To enable transmission to one of the routers 14, 16 and 18 that lie outside the first network, the Ethernet network 22 transmits each Ethernet-formatted information frame 20 destined for one of the routers 14, 16 and 18 to an Ethernet Internet Working Switch (EIWS) 24 for transmission to a Wide Area ATM core network 26 (a second) network that serves the routers 14, 16 and 18 as discussed below. The EIWS 24 typically comprises a Ethernet switch that serves as a proxy between the Ethernet network 22 and the ATM network 26 which services a plurality of edge devices 28, 30, and 32 that utilize one of a plurality of protocols, such as ATM or Frame Relay. Each of the edge devices 28, 30 and 32 forwards traffic between ATM core network 26 and one of the destination routers 14, 16, and 18, respectively, across one of PVCs 34, 36, and 38, respectively.

[0012] FIGURE 2 illustrates the manner in which EIWS 24 acts as an Ethernet Address Resolution Protocol (ARP) proxy for the ATM domain that serves the edge routers 28, 30 and 32 to initially resolve the protocol disparity between the networks 22 and 26 to subsequently permit transmission of frames from the router 12 to one of the routers 14, 16 and 18 (all of FIG. 1) and vice versa. To resolve the protocol disparity, the router 12 first floods an ARP polling request in the form of a packet to determine which destinations are on the Ethernet network 22 directly as well as which destinations (e.g., the routers 14, 16 and 18 of FIG. 1) are available through the ATM network 26. The EIWS 24 receives the ARP polling request and opens the packet to determine the Virtual Local Area Network (VLAN) tag embodied in the packet for the purpose of matching the VLAN tag to a path (i.e., a PVC) in the ATM network 26. To match the VLAN tag to an ATM PVC, the EIWS 24 uses a VLAN to PVC mapping table. The EIWS 24 then sends an ATM-encoded packet out the PVC through the ATM network 26 to the destination ATM router 28. The ATM router 28 receives the polling request and responds back across the same PVC to the ATM network 26 with the IP address of the router formatted as an IP packet. The ATM network 26 forwards this packet to the EIWS 24, which then opens the packet, and obtains the IP Address for the remote ATM edge device 28. The EIWS 24 then encodes an ARP reply for the source with an Ethernet Source MAC Address from a locally available pool of Addresses. The ARP reply also contains the IP

Address from ATM router 28 of the destination router 14. Furthermore, the ARP reply is encoded with same VLAN that came from router 12. Router 12 receives the ARP reply and updates its ARP table. At this point, the router 12 believes that ATM router 28 is directly connected to its own Ethernet port. However, in actuality, the EIWS 24 acts as a proxy in this exchange. The router 12 would follow the same method to resolve the address of the edge devices 30 and 32 and the routers 16 and 18, respectively, served thereby. Having resolved the addresses, the EIWS 24 can then facilitate the actual transmission of data from the router 12 to one of the routes 14, 16 and 18 as described below. Although not described, each of the ATM routers 28, 30 and 32 would resolve the destination address for information frames sent to the router 12 in a comparable manner.

[0013] FIGURE 3 best illustrates the manner in which the EIWS 24 operates to interwork the Ethernet-based MAN 22 with the ATM network 26 (both of FIG. 1) to facilitate the actual transmission of data following the initial address resolution discussed above with respect to FIG. 2. The EIWS 24 enables the transmission of information embodied in the payload of each of a plurality of Ethernet-formatted information frames, illustratively represented by frames 20₁ and 20₂, by forming corresponding ATM cell sequences, illustratively represented by ATM cell sequences 40₁ and 40₂ respectively. Each Ethernet formatted information frame received at the EIWS 24, such as frame 20₁ includes an Ethernet header 42, a payload 44 and a VLAN tag 46. The Ethernet header 42 contains certain administrative data, such as the identity of the source of the frame (the Ethernet source MAC Address) and the identity of the destination (the Ethernet destination MAC Address). The payload 44 contains the information of interest (e.g. a IP packet), whereas the VLAN tag 46 contains the endpoint destination address, in the form of a sub-interface on the router 12 associated with a corresponding one of the destination routers 14, 16, and 18 of FIG. 1. The Ethernet destination MAC Address must be for the EIWS 24 and the EIWS must act as a proxy for the ATM portion of the end-to-end interconnection. The VLAN tag provides the ability to reach the ATM endpoint. Unfortunately, the address embodied in the VLAN tag 46 has no meaning to the ATM network 26 and thus, the ATM network could not by itself readily forward the frame 20₁ to one of the destination routers 14, 16, and 18.

[0014] In accordance with present principles, the EIWS 24, upon receipt of an Ethernet formatted frame, such as frame 20₁, first determines whether the frame is destined for an endpoint served by the ATM network 26, such as one of the routers 28, 30, and 32 of FIG. 1, based on the address specified by the VLAN tag 46 of FIG. 2 in that Ethernet-formatted frame. If the frame 20₁ is indeed destined for such an endpoint, then the frame requires interworking, whereupon, the EIWS 24 then removes the both Ethernet header 42 and the VLAN tag 46 from the frame, leaving just the Ethernet payload 44 which is nothing more than an IP Packet. The EIWS 24 then forms ATM AAL5 Frame 44₁ that includes this payload (the IP packet).

[0015] The EIWS 24 then determines the destination address for the AAL5 Frame 44₁ (i.e., the appropriate ATM PVC, such as one of the PVCs 34, 36 and 38 of FIG. 1 serving the routers 14, 16 and 18, respectively) by mapping the VLAN tag to a corresponding PVC via a mapping table 46 that cross-references VLAN addresses to corresponding ATM PVCs. By mapping the VLAN tag 46 to the corresponding ATM PVC, the EIWS 24 effectively converts the Ethernet address into an ATM address. The EIWS 24 then breaks down the AAL5 frame 42 in a corresponding sequence of ATM cells, such as cell sequence 40₁, for transmission to the ATM network 26. In a similar fashion, the EIWS 24 will remove the Ethernet header 42 and VLAN tag 46 from a subsequent Ethernet-formatted frame 20₂ and encapsulate its payload 44 into another AAL5 frame 44₂. The EIWS 24 then maps the VLAN tag of the frame 20₂ to a corresponding PVC to yield a subsequent ATM cell sequence 40₂.

[0016] The ATM network 26 receives the ATM cell sequences 40₁ and 40₂ and transmits them to the appropriate destination (i.e., the corresponding one of endpoint routers 14, 16 and 18 served by the edge devices 28, 30 and 32, respectively), based on the PVC value obtained from the mapping performed by the EIWS 24. Upon receipt of an ATM cell sequence, the endpoint router, e.g., endpoint router 14, knows that the cell sequence contains an AAL5 payload. After removing the administrative information from the ATM cell sequence, the endpoint router is left with the payload in the form of an IP Packet. At that time, the endpoint router (e.g. router 14) makes a *routing* decision based on it's own routing table. In the illustrative embodiment of FIG. 1, the endpoint routers 14, 16, and 18 all believe that they are connected to an ATM network that serves

the source router 12 rather than Ethernet MAN 22 which actually serves the source router.

[0017] In addition to interworking Ethernet frames to ATM frames, the EIWS 24 also interworks ATM frames to Ethernet frames. One or more of the endpoint routers 14, 16, and 18 may originate a sequence of ATM cells that embody a payload containing information of interest that is ultimately destined for the router 12.. Such ATM cell sequences originated by one of the routers 14-18 will include a VPN PVC that serves as the ATM address for the router 12. Thus, the sending router (e.g., router 14) perceives the receiving router 12 as an ATM router despite its actual status as an Ethernet router. The sequence of ATM cells from the sending router 14 passes to the corresponding edge router (e.g., router 28) for forwarding to the ATM core network 26 and subsequent transmission to the EIWS 24.

[0018] The EIWS 24 receives the ATM cell sequences and translates them into usable Ethernet frames by essentially reversing the process previously. The EIWS 24 receives each ATM cell sequence and then maps the VPN PVC information therein to a corresponding VLAN tag using the table 46. The EIWS 24 then combines each cell sequence to reassemble the AAL5 payload and thereafter strips out the AAL5 administrative information to yield the remaining payload, which is inserted into an Ethernet Frame with the specified VLAN tag. The EIWS 24 then forwards the Ethernet frame (including the appropriate VLAN tag 26) to the Ethernet-based MAN for transmission to the router 12. When constructing the Ethernet Frame, the MAC Address of Router 12 must be inserted in the Ethernet Destination MAC Address field. The Source MAC Address must be provided as a proxy function by the EIWS as depicted in FIG. 2.

[0019] Router 12 receives the VLAN-tagged Ethernet frame and strips the Ethernet header 42 and the VLAN tag 46, leaving the payload 44 comprising an IP Packet. Router 12 then makes a *routing* decision based on its own routing table. In this way Router 12 receives an Ethernet Frame that it believes were originated from an Ethernet-based router, notwithstanding the fact that the sending Router 14 for example is actually connected via an ATM or Frame relay link to a corresponding edge router 28.

2001-0381 - Holmgren-Kinsky-Medamanna-Szela
November 30, 2001

[0020] The foregoing describes a technique for interworking different types of data networks that have protocols and different addressing schemes.

[0021] The above-described embodiments merely illustrate the principles of the invention. Those skilled in the art may make various modifications and changes that will embody the principles of the invention and fall within the spirit and scope thereof.